

FIG. 1

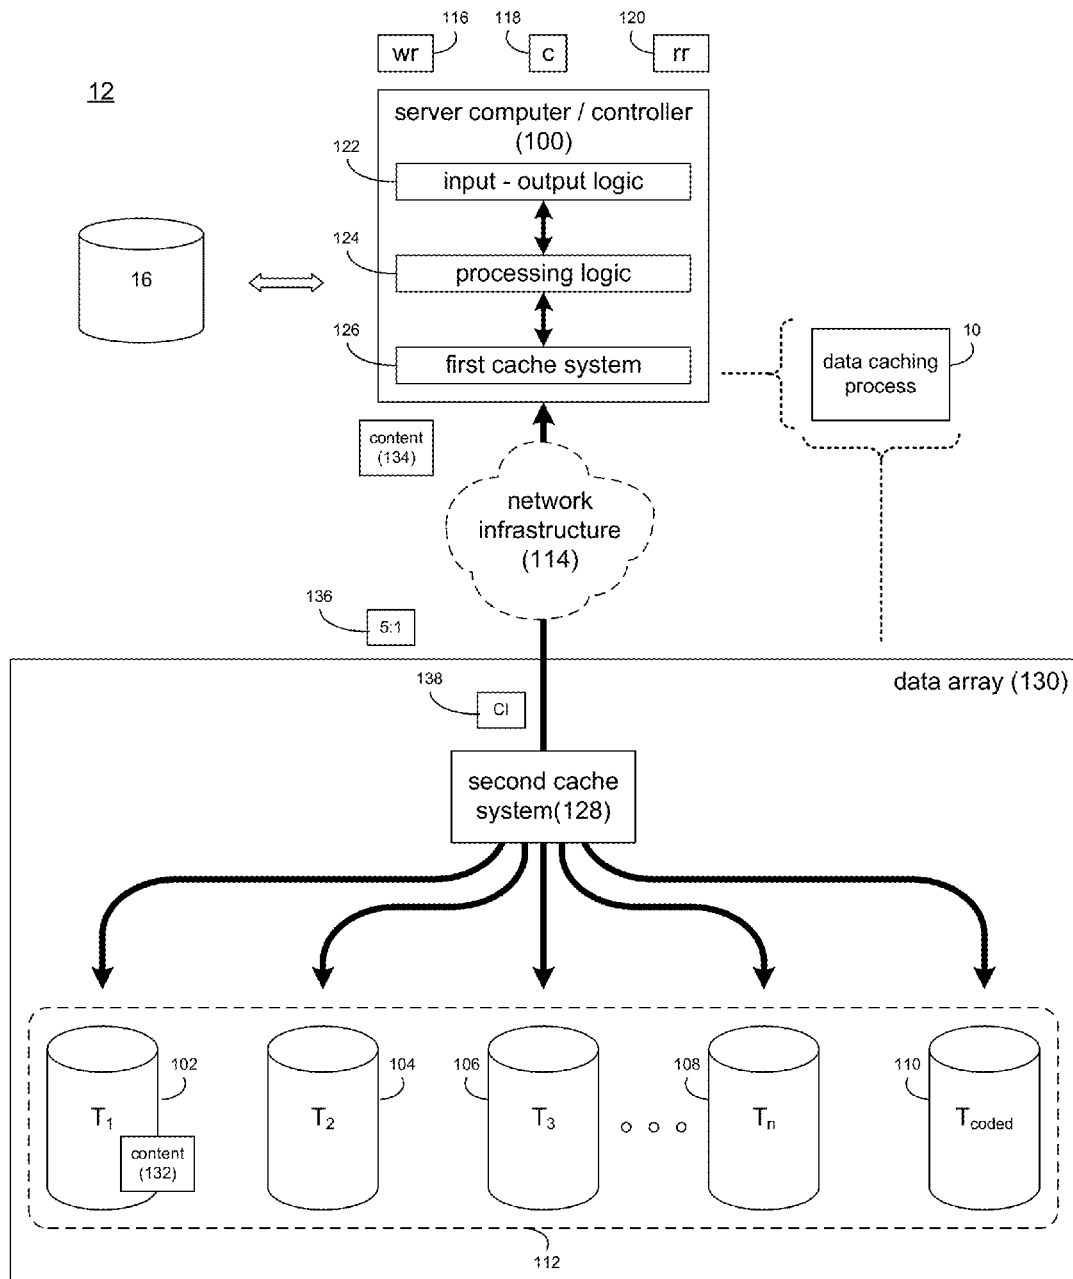


FIG. 2



FIG. 4




FIG. 3




FIG. 5

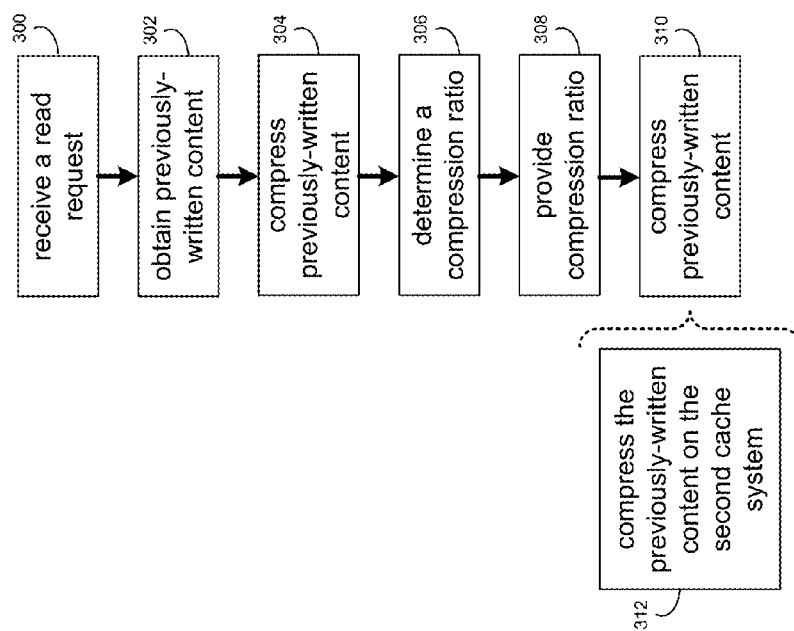


FIG. 6

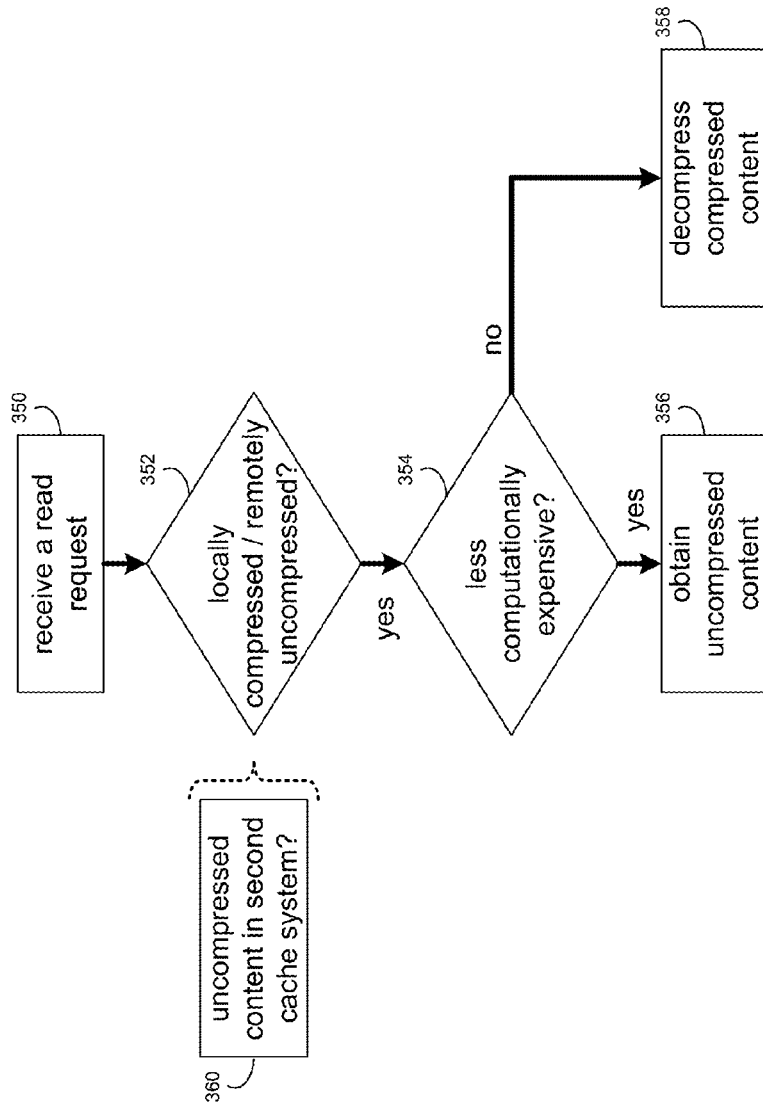


FIG. 7

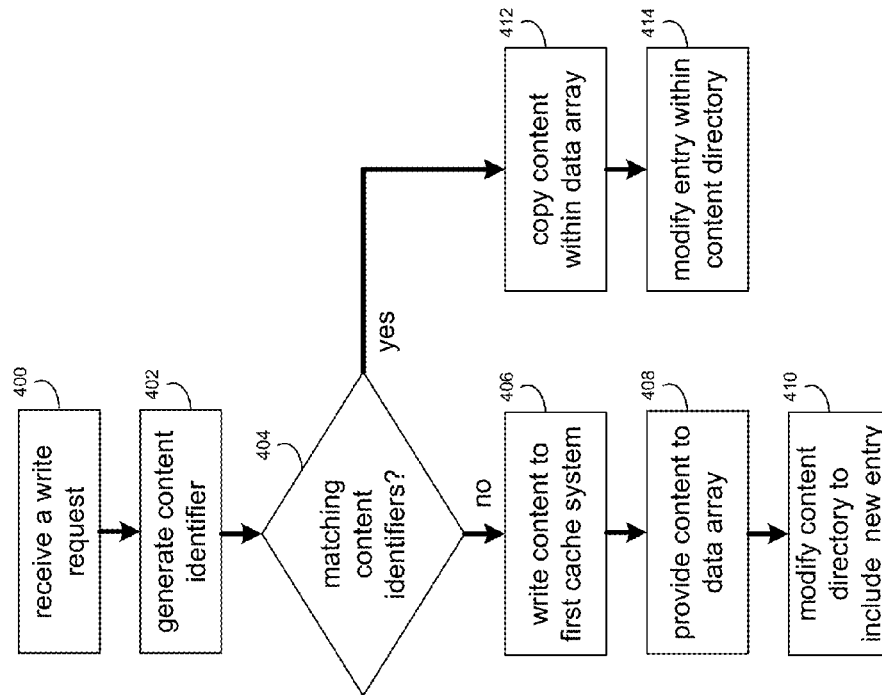


FIG. 8

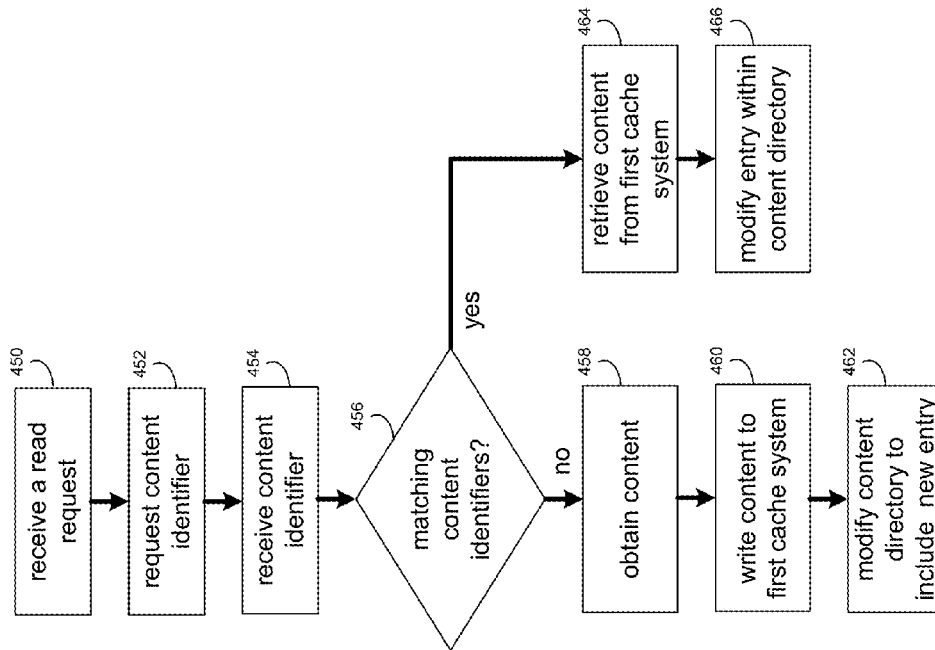


FIG. 9

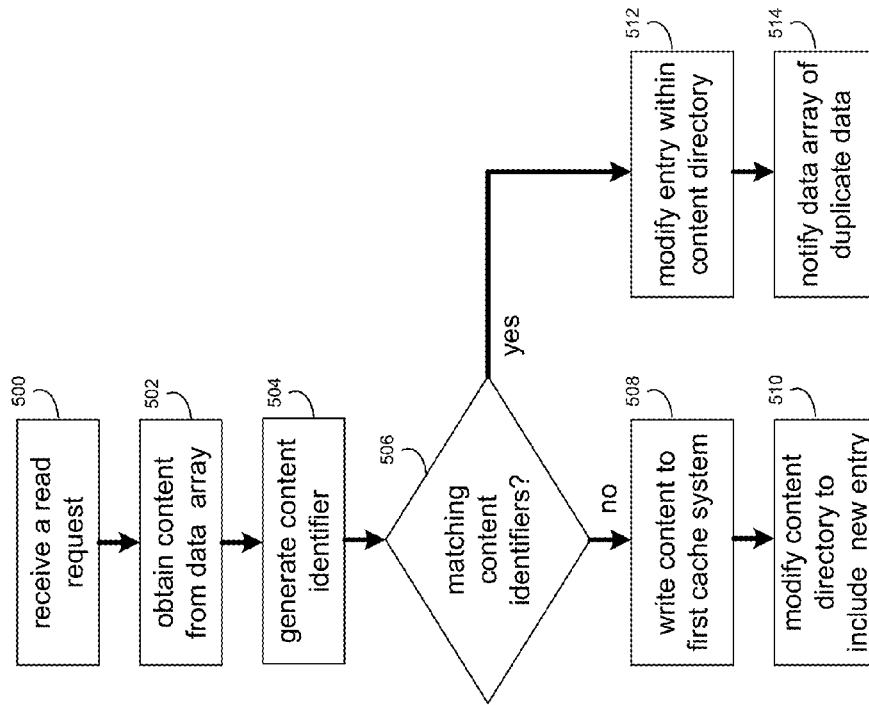


FIG. 10

## SYSTEM AND METHOD FOR IMPROVING CACHE PERFORMANCE

### RELATED APPLICATIONS

The subject application is a continuation application of U.S. patent application with Ser. No. 13/308,345, filed on Nov. 30, 2011, the entire content of which is herein incorporated by reference.

### TECHNICAL FIELD

This disclosure relates to cache memory systems and, more particularly, to systems and methods for improving the performance of cache memory systems.

### BACKGROUND

Storing and safeguarding electronic content is of paramount importance in modern business. Accordingly, various systems may be employed to protect such electronic content.

The use of solid-state storage devices is increasing in popularity. A solid state storage device is a content storage device that uses solid-state memory to store persistent content. A solid-state storage device may emulate (and therefore replace) a conventional hard disk drive. Additionally/alternatively, a solid state storage device may be used within a cache memory system. With no moving parts, a solid-state storage device largely eliminates (or greatly reduces) seek time, latency and other electromechanical delays and failures associated with a conventional hard disk drive.

### SUMMARY OF DISCLOSURE

In a first implementation, a computer-implemented method of processing read requests includes receiving a read request on a first cache system, wherein the read request identifies previously-written content included within a data array. The previously-written content identified in the read request is obtained from the data array. A read request content identifier is generated for the previously-written content identified in the read request. The read request content identifier associated with the previously-written content identified in the read request is compared to a plurality of content identifiers included within a content directory for the first cache system to determine if a matching content identifier exists. Each of the plurality of content identifiers is associated with a piece of previously-written content included within the first cache system.

One or more of the following features may be included. If a matching content identifier is not identified, the previously-written content identified in the read request may be written to the first cache system. If a matching content identifier is identified, the data array may be notified that the previously-written content identified in the read request is identical to previously-written content associated with the matching content identifier. The read request may identify a logical unit and a storage address of the previously-written content identified in the read request on the data array.

If a matching content identifier is identified, an entry within the content directory that is associated with the matching content identifier may be modified to identify the logical storage unit and the storage address of the previously-written content identified in the read request on the data array. The read request content identifier may be a hash function of the previously-written content identified in the read request. The content directory for the first cache system may include a plurality of entries, wherein each of the plurality of entries is associated with a specific piece of previously-written content within the first cache system. Each of the plurality of entries

within the content directory may identify: a logical storage unit and a storage address at which the specific piece of previously-written content is located within the storage array; a first cache address at which the specific piece of previously-written content is located within the first cache system, and a content identifier for the specific piece of previously-written content.

In another implementation, a computer program product resides on a computer readable medium that has a plurality of instructions stored on it. When executed by a processor, the instructions cause the processor to perform operations including receiving a read request on a first cache system, wherein the read request identifies previously-written content included within a data array. The previously-written content identified in the read request is obtained from the data array. A read request content identifier is generated for the previously-written content identified in the read request. The read request content identifier associated with the previously-written content identified in the read request is compared to a plurality of content identifiers included within a content directory for the first cache system to determine if a matching content identifier exists. Each of the plurality of content identifiers is associated with a piece of previously-written content included within the first cache system.

One or more of the following features may be included. If a matching content identifier is not identified, the previously-written content identified in the read request may be written to the first cache system. If a matching content identifier is identified, the data array may be notified that the previously-written content identified in the read request is identical to previously-written content associated with the matching content identifier. The read request may identify a logical unit and a storage address of the previously-written content identified in the read request on the data array.

If a matching content identifier is identified, an entry within the content directory that is associated with the matching content identifier may be modified to identify the logical storage unit and the storage address of the previously-written content identified in the read request on the data array. The read request content identifier may be a hash function of the previously-written content identified in the read request. The content directory for the first cache system may include a plurality of entries, wherein each of the plurality of entries is associated with a specific piece of previously-written content within the first cache system. Each of the plurality of entries within the content directory may identify: a logical storage unit and a storage address at which the specific piece of previously-written content is located within the storage array; a first cache address at which the specific piece of previously-written content is located within the first cache system, and a content identifier for the specific piece of previously-written content.

In another implementation, a computing system includes at least one processor and at least one memory architecture coupled with the at least one processor, wherein the computing system is configured to perform operations including receiving a read request on a first cache system, wherein the read request identifies previously-written content included within a data array. The previously-written content identified in the read request is obtained from the data array. A read request content identifier is generated for the previously-written content identified in the read request. The read request content identifier associated with the previously-written content identified in the read request is compared to a plurality of content identifiers included within a content directory for the first cache system to determine if a matching content identifier exists. Each of the plurality of content identifiers is associated with a piece of previously-written content included within the first cache system.

One or more of the following features may be included. If a matching content identifier is not identified, the previously-written content identified in the read request may be written to the first cache system. If a matching content identifier is identified, the data array may be notified that the previously-written content identified in the read request is identical to previously-written content associated with the matching content identifier. The read request may identify a logical unit and a storage address of the previously-written content identified in the read request on the data array.

If a matching content identifier is identified, an entry within the content directory that is associated with the matching content identifier may be modified to identify the logical storage unit and the storage address of the previously-written content identified in the read request on the data array. The read request content identifier may be a hash function of the previously-written content identified in the read request. The content directory for the first cache system may include a plurality of entries, wherein each of the plurality of entries is associated with a specific piece of previously-written content within the first cache system. Each of the plurality of entries within the content directory may identify: a logical storage unit and a storage address at which the specific piece of previously-written content is located within the storage array; a first cache address at which the specific piece of previously-written content is located within the first cache system, and a content identifier for the specific piece of previously-written content.

The details of one or more implementations are set forth in the accompanying drawings and the description below. Other features and advantages will become apparent from the description, the drawings, and the claims.

#### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagrammatic view of a storage system and a data caching process coupled to a distributed computing network;

FIG. 2 is a diagrammatic view of the storage system of FIG. 1;

FIG. 3 is a diagrammatic view of a data write request for use with the data caching process of FIG. 1;

FIG. 4 is a diagrammatic view of a data read request for use with the data caching process of FIG. 1;

FIG. 5 is a diagrammatic view of a content directory for use with the data caching process of FIG. 1;

FIG. 6 is a first flow chart of the data caching process of FIG. 1;

FIG. 7 is a second flow chart of the data caching process of FIG. 1;

FIG. 8 is a third flow chart of the data caching process of FIG. 1;

FIG. 9 is a fourth flow chart of the data caching process of FIG. 1; and

FIG. 10 is a fifth flow chart of the data caching process of FIG. 1.

Like reference symbols in the various drawings indicate like elements.

#### DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

##### System Overview:

As will be appreciated by one skilled in the art, the present disclosure may be embodied as a method, system, or computer program product. Accordingly, the present disclosure may take the form of an entirely hardware embodiment, an

entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, the present disclosure may take the form of a computer program product on a computer-usable storage medium having computer-usable program code embodied in the medium.

Any suitable computer usable or computer readable medium may be utilized. The computer-usable or computer-readable medium may be, for example but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, device, or propagation medium. More specific examples (a non-exhaustive list) of the computer-readable medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a transmission media such as those supporting the Internet or an intranet, or a magnetic storage device. Note that the computer-usable or computer-readable medium could even be paper or another suitable medium upon which the program is printed, as the program can be electronically captured, via, for instance, optical scanning of the paper or other medium, then compiled, interpreted, or otherwise processed in a suitable manner, if necessary, and then stored in a computer memory. In the context of this document, a computer-usable or computer-readable medium may be any medium that can contain, store, communicate, propagate, or transport the program for use by or in connection with the instruction execution system, apparatus, or device. The computer-usable medium may include a propagated data signal with the computer-usable program code embodied therewith, either in baseband or as part of a carrier wave. The computer usable program code may be transmitted using any appropriate medium, including but not limited to the Internet, wireline, optical fiber cable, RF, etc.

Computer program code for carrying out operations of the present disclosure may be written in an object oriented programming language such as Java, Smalltalk, C++ or the like. However, the computer program code for carrying out operations of the present disclosure may also be written in conventional procedural programming languages, such as the "C" programming language or similar programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

The present disclosure is described below with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the disclosure. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer

or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

These computer program instructions may also be stored in a computer-readable memory that can direct a computer or other programmable data processing apparatus to function in a particular manner, such that the instructions stored in the computer-readable memory produce an article of manufacture including instruction means which implement the function/act specified in the flowchart and/or block diagram block or blocks.

The computer program instructions may also be loaded onto a computer or other programmable data processing apparatus to cause a series of operational steps to be performed on the computer or other programmable apparatus to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide steps for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

System Overview:

Referring to FIG. 1, there is shown data caching process 10 that may reside on and may be executed by storage system 12, which may be connected to network 14 (e.g., the Internet or a local area network). Examples of storage system 12 may include, but are not limited to: a Network Attached Storage (NAS) system, a Storage Area Network (SAN), a personal computer with a memory system, a server computer with a memory system, and a cloud-based device with a memory system.

As is known in the art, a SAN may include one or more of a personal computer, a server computer, a series of server computers, a mini computer, a mainframe computer, a RAID device and a NAS system. The various components of storage system 12 may execute one or more operating systems, examples of which may include but are not limited to: Microsoft Windows XP Server™; Novell Netware Redhat Linux™, Unix, or a custom operating system, for example.

The instruction sets and subroutines of data caching process 10, which may be stored on storage device 16 included within storage system 12, may be executed by one or more processors (not shown) and one or more memory architectures (not shown) included within storage system 12. Storage device 16 may include but is not limited to: a hard disk drive; a tape drive; an optical drive; a RAID device; a random access memory (RAM); a read-only memory (ROM); and all forms of flash memory storage devices.

Network 14 may be connected to one or more secondary networks (e.g., network 18), examples of which may include but are not limited to: a local area network; a wide area network; or an intranet, for example.

Various data requests (e.g. data request 20) may be sent from client applications 22, 24, 26, 28 to storage system 12. Examples of data request 20 may include but are not limited to data write requests (i.e. a request that content be written to storage system 12) and data read requests (i.e. a request that content be read from storage system 12).

The instruction sets and subroutines of client applications 22, 24, 26, 28, which may be stored on storage devices 30, 32, 34, 36 (respectively) coupled to client electronic devices 38, 40, 42, 44 (respectively), may be executed by one or more processors (not shown) and one or more memory architectures (not shown) incorporated into client electronic devices 38, 40, 42, 44 (respectively). Storage devices 30, 32, 34, 36 may include but are not limited to: hard disk drives; tape drives; optical drives; RAID devices; random access memories (RAM); read-only memories (ROM), and all forms of

flash memory storage devices. Examples of client electronic devices 38, 40, 42, 44 may include, but are not limited to, personal computer 38, laptop computer 40, personal digital assistant 42, notebook computer 44, a server (not shown), a data-enabled, cellular telephone (not shown), and a dedicated network device (not shown).

Users 46, 48, 50, 52 may access storage system 12 directly through network 14 or through secondary network 18. Further, storage system 12 may be connected to network 14 through secondary network 18, as illustrated with link line 54.

The various client electronic devices may be directly or indirectly coupled to network 14 (or network 18). For example, personal computer 38 is shown directly coupled to network 14 via a hardwired network connection. Further, notebook computer 44 is shown directly coupled to network 18 via a hardwired network connection. Laptop computer 40 is shown wirelessly coupled to network 14 via wireless communication channel 56 established between laptop computer 40 and wireless access point (i.e., WAP) 58, which is shown directly coupled to network 14. WAP 58 may be, for example, an IEEE 802.11a, 802.11b, 802.11g, 802.11n, Wi-Fi, and/or Bluetooth device that is capable of establishing wireless communication channel 56 between laptop computer 40 and WAP 58. Personal digital assistant 42 is shown wirelessly coupled to network 14 via wireless communication channel 60 established between personal digital assistant 42 and cellular network/bridge 62, which is shown directly coupled to network 14.

As is known in the art, all of the IEEE 802.11x specifications may use Ethernet protocol and carrier sense multiple access with collision avoidance (i.e., CSMA/CA) for path sharing. The various 802.11x specifications may use phase-shift keying (i.e., PSK) modulation or complementary code keying (i.e., CCK) modulation, for example. As is known in the art, Bluetooth is a telecommunications industry specification that allows e.g., mobile phones, computers, and personal digital assistants to be interconnected using a short-range wireless connection.

Client electronic devices 38, 40, 42, 44 may each execute an operating system, examples of which may include but are not limited to Microsoft Windows™, Microsoft Windows CE™, Redhat Linux™, or a custom operating system.

The Data Caching Process:

For the following discussion, client application 22 is going to be described for illustrative purposes. However, this is not intended to be a limitation of this disclosure, as other client applications (e.g., client applications 24, 26, 28) may be equally utilized.

For illustrative purposes, storage system 12 will be described as being a network-based storage system that includes a plurality of electro-mechanical backend storage devices. However, this is for illustrative purposes only and is not intended to be a limitation of this disclosure, as other configurations are possible and are considered to be within the scope of this disclosure. For example and as discussed above, storage system 12 may be a personal computer that includes a single electro-mechanical storage device.

Referring also to FIG. 2, storage system 12 may include a server computer/controller (e.g. server computer/controller 100), and a plurality of storage targets  $T_{1..n}$  (e.g. storage targets 102, 104, 106, 108). Storage targets 102, 104, 106, 108 may be configured to provide various levels of performance and/or high availability. For example, one or more of storage targets 102, 104, 106, 108 may be configured as a RAID 0 array, in which data is striped across storage targets. By striping data across a plurality of storage targets, improved performance may be realized. However, RAID 0 arrays do not

provide a level of high availability. Accordingly, one or more of storage targets **102**, **104**, **106**, **108** may be configured as a RAID 1 array, in which data is mirrored between storage targets. By mirroring data between storage targets, a level of high availability is achieved as multiple copies of the data are stored within storage system **12**.

While storage targets **102**, **104**, **106**, **108** are discussed above as being configured in a RAID 0 or RAID 1 array, this is for illustrative purposes only and is not intended to be a limitation of this disclosure, as other configurations are possible. For example, storage targets **102**, **104**, **106**, **108** may be configured as a RAID 3, RAID 4, RAID 5 or RAID 6 array.

While in this particular example, storage system **12** is shown to include four storage targets (e.g. storage targets **102**, **104**, **106**, **108**), this is for illustrative purposes only and is not intended to be a limitation of this disclosure. Specifically, the actual number of storage targets may be increased or decreased depending upon e.g. the level of redundancy/performance/capacity required.

Storage system **12** may also include one or more coded targets **110**. As is known in the art, a coded target may be used to store coded data that may allow for the regeneration of data lost/corrupted on one or more of storage targets **102**, **104**, **106**, **108**. An example of such a coded target may include but is not limited to a hard disk drive that is used to store parity data within a RAID array.

While in this particular example, storage system **12** is shown to include one coded target (e.g., coded target **110**), this is for illustrative purposes only and is not intended to be a limitation of this disclosure. Specifically, the actual number of coded targets may be increased or decreased depending upon e.g. the level of redundancy/performance/capacity required.

Examples of storage targets **102**, **104**, **106**, **108** and coded target **110** may include one or more electro-mechanical hard disk drives, wherein a combination of storage targets **102**, **104**, **106**, **108** and coded target **110** may form non-volatile, electro-mechanical memory system **112**.

The manner in which storage system **12** is implemented may vary depending upon e.g. the level of redundancy/performance/capacity required. For example, storage system **12** may be a RAID device in which server computer/controller **100** is a RAID controller card and storage targets **102**, **104**, **106**, **108** and/or coded target **110** are individual "hot-swappable" hard disk drives. An example of such a RAID device may include but is not limited to an NAS device. Alternatively, storage system **12** may be configured as a SAN, in which server computer/controller **100** may be e.g., a server computer and each of storage targets **102**, **104**, **106**, **108** and/or coded target **110** may be a RAID device and/or computer-based hard disk drive. Further still, one or more of storage targets **102**, **104**, **106**, **108** and/or coded target **110** may be a SAN.

In the event that storage system **12** is configured as a SAN, the various components of storage system **12** (e.g. server computer/controller **100**, storage targets **102**, **104**, **106**, **108**, and coded target **110**) may be coupled using network infrastructure **114**, examples of which may include but are not limited to an Ethernet (e.g., Layer 2 or Layer 3) network, a fiber channel network, an InfiniBand network, or any other circuit switched/packet switched network.

Storage system **12** may execute all or a portion of data caching process **10**. The instruction sets and subroutines of data caching process **10**, which may be stored on a storage device (e.g., storage device **16**) coupled to server computer/controller **100**, may be executed by one or more processors (not shown) and one or more memory architectures (not

shown) included within server computer/controller **100**. Storage device **16** may include but is not limited to: a hard disk drive; a tape drive; an optical drive; a RAID device; a random access memory (RAM); a read-only memory (ROM); and all forms of flash memory storage devices.

As discussed above, various data requests (e.g. data request **20**) may be generated. For example, these data requests may be sent from client applications **22**, **24**, **26**, **28** to storage system **12**. Additionally/alternatively and when server computer/controller **100** is configured as an application server, these data requests may be internally generated within server computer/controller **100**. Examples of data request **20** may include but are not limited to data write request **116** (i.e. a request that content **118** be written to storage system **12**) and data read request **120** (i.e. a request that content **118** be read from storage system **12**).

Server computer/controller **100** may include input-output logic **122** (e.g., a network interface card), processing logic **124**, and first cache system **126**. Examples of first cache system **126** may include but are not limited to a volatile, solid-state, cache memory system (e.g., a dynamic RAM cache memory system) and/or a non-volatile, solid-state, cache memory system (e.g., a flash-based, cache memory system).

During operation of server computer/controller **100**, content **118** to be written to storage system **12** may be received by input-output logic **122** (e.g. from network **14** and/or network **18**) and processed by processing logic **124**. Additionally/alternatively and when server computer/controller **100** is configured as an application server, content **118** to be written to storage system **12** may be internally generated by server computer/controller **100**. As will be discussed below in greater detail, processing logic **124** may initially store content **118** within first cache system **126**.

Depending on the manner in which first cache system **126** is configured, processing logic **124** may immediately write content **118** to second cache system **128**/non-volatile, electro-mechanical memory system **112** (if first cache system **126** is configured as a write-through cache) or may subsequently destage content **118** to second cache system **128**/non-volatile, electro-mechanical memory system **112** (if first cache system **126** is configured as a write-back cache). Additionally and in certain configurations, processing logic **124** may calculate and store coded data on coded target **110** (included within non-volatile, electromechanical memory system **112**) that may allow for the regeneration of data lost/corrupted on one or more of storage targets **102**, **104**, **106**, **108**. For example, if processing logic **124** was included within a RAID controller card or a NAS/SAN controller, processing logic **124** may calculate and store coded data on coded target **110**. However, if processing logic **124** was included within e.g., an applications server, data array **130** may calculate and store coded data on coded target **110**.

The combination of second cache system **128** and non-volatile, electromechanical memory system **112** may form data array **130**, wherein first cache system **126** may be sized so that the number of times that data array **130** is accessed may be reduced. Accordingly, by sizing first cache system **126** so that first cache system **126** retains a quantity of data sufficient to satisfy a significant quantity of data requests (e.g., data request **20**), the overall performance of storage system **12** may be enhanced. As will be described below in greater detail, first cache system **126** may be a content-aware cache system.

Further, second cache system **128** within data array **130** may be sized so that the number of times that non-volatile, electromechanical memory system **112** is accessed may be

reduced. Accordingly, by sizing second cache system **128** so that second cache system **128** retains a quantity of data sufficient to satisfy a significant quantity of data requests (e.g., data request **20**), the overall performance of storage system **12** may be enhanced. As will be described below in greater detail, second cache system **130** may be a content-aware cache system.

As discussed above, the instruction sets and subroutines of data caching process **10**, which may be stored on storage device **16** included within storage system **12**, may be executed by one or more processors (not shown) and one or more memory architectures (not shown) included within storage system **12**. Accordingly, in addition to being executed on server computer/controller **100**, some or all of the instruction sets and subroutines of data caching process **10** may be executed by one or more processors (not shown) and one or more memory architectures (not shown) included within data array **130**.

Referring also to FIGS. 3-4, data request **20** (e.g. data read request **116** and/or data write request **120**) may be processed by server computer/controller **100** to extract pertinent information concerning these data requests.

When data request **20** is a data write request (e.g., write request **116**), write request **116** may include content **118** to be written to data array **130**. Additionally, write request **116** may include a storage address **200** that defines the intended storage location within storage array **130** at which content **118** is to be stored. For example, storage address **200** may define a particular logical unit within data array **130** (e.g., a LUN or Logical Unit Number) and a particular storage address within that specific logical unit (e.g., an LBA or Logical Block Address) for storing content **118**.

Concerning read request **120**, these requests do not include any content to be written to data array **130**, as these are read requests and concern content to be read from data array **130**. Read request **120** may include a storage address **202** that defines the storage location within storage array **130** from which content is to be retrieved. For example, storage address **202** may define a particular logical unit within data array **130** (e.g., a LUN or Logical Unit Number) and a particular storage address within that specific logical unit (e.g., an LBA or Logical Block Address) for retrieving the content sought from data array **130**.

As will be discussed below in greater detail and referring also to FIG. 5, data caching process **10** may maintain content directory **250**, which may be used to locate various pieces of content within first cache system **126**. In one particular embodiment of content directory **250**, content directory **250** may include plurality of entries **252**, wherein each of these entries may identify: data array storage address **200/202** (e.g. a logical storage unit and a storage address at which a specific piece of previously-written content is located within storage array **130**); first cache address **254** (e.g., the location within first cache system **126** at which the specific piece of previously-written content is also located), and content identifier **256** for the specific piece of previously-written content. Accordingly, content directory **250** may identify the location of specific pieces of content included within first cache system **126** and their corresponding pieces of data within data array **130**, as well as a content identifier that uniquely identifies the specific piece of content.

Content identifier **256** may be a mathematical representation of the specific piece of previously-written content that may allow e.g. server computer/controller **100** to quickly determine whether two pieces of previously-written content are identical, as identical pieces of content would have identical content identifiers. In one particular embodiment, con-

tent identifier **256** may be a hash function (e.g., a cryptographic hash) of the previously-written content.

Compression Ratio Monitoring:

Data caching process **10** may be configured to compress the data stored on data array **130** to conserve storage space if computationally desirable.

Referring also to FIG. 6, data caching process **10** may be configured to receive **300** read request **120** on first cache system **126**, wherein read request **120** identifies previously-written content (as defined by storage address **202**) included within data array **130**. For example, assume that user **46** is using client application **22** to access data (i.e. content **132**) that is currently being stored on data array **130**. Accordingly, client application **22** may generate read request **120** which, as discussed above, may define a particular logical unit within data array **130** (e.g., a LUN or Logical Unit Number) and a particular storage address within that specific logical unit (e.g., an LBA or Logical Block Address) for retrieving content **132** sought from data array **130** by client application **22**.

Assume that read request **120** defines LUN0/LBA5 as the location of content **132** within data array **130**. Upon receiving **300** read request **120**, data caching process **10** may compare the location of content **132** within data array **130** (namely LUN0/LBA5) with each of the plurality of entries **252** defined within content directory **250** to determine if a copy of content **132** is locally available (i.e., cached) within first cache system **126**. If LUN0/LBA5 was defined within content directory **250** (meaning that a local cached copy of content **132** is present/available within first cache system **126**), that particular entry would also define a corresponding first cache address (e.g. first cache address **254**) within first cache system **126** at which content **132** would be locally-available and retrievable from the first cache system **126**. Conversely, in the event that LUN0/LBA5 is not defined within content directory **250** (meaning that a local cached copy of content **132** is not present/available within first cache system **126**), data caching process **10** may need to obtain **302** content **132** identified in read request **120** from data array **130**. In this particular example, since LUN0/LBA5 is not defined within content directory **250**, a local cached copy of content **132** is not present/available within first cache system **126** and data caching process **10** will be need to obtain **302** content **132** from data array **130**.

Once content **132** is obtained **302** by data caching process **10**, data caching process **10** may compress **304** content **132** prior to content **132** being stored within first cache system **126**, thus generating compressed content **134** which may be stored within (and conserve the storage space of) first cache system **126**. When compressing **304** content **132** to generate compressed content **134**, data caching process **10** may determine **306** compression ratio **136** for compressed content **134**. For example, if content **132** is 1.00 MB in size and compressed content **134** is 200 kB in size, data caching process **10** may determine **306** compression ratio **136** of 5:1 for compressed content **134**. Data caching process **10** may provide **308** compression ratio **136** to data array **130**.

Upon receiving compression ratio **136**, data caching process **10** may compare compression ratio **136** to a predefined threshold to determine whether content **132** (stored on data array **130**) should be compressed **310**. For example, assume for illustrative purposes that this predefined threshold is 3:1, wherein if the compression ratio is less than 3:1, content **132** (stored on data array **130**) would not be compressed **310** by data caching process **10**, as the computational expense of compressing content **132** (for storage on data array **130**) would not be offset by the space savings achieved. However, if the compression ratio is greater than or equal to 3:1, data

11

caching process 10 would compress 310 content 132, as the space savings achieved on data array 130 would offset the computational expense of compressing content 132.

As discussed above, data array 130 may include second cache system 128. Accordingly, if data caching process 10 compresses 310 content 132 on data array 130, if data array 130 also includes second cache system 128, compressing 310 content 132 may also include/may alternatively include compressing 312 any cached copies of content 132 stored on second cache system 128.

Uncompressed Data Retrieval:

Data caching process 10 may be configured to minimize/reduce processor loading by intelligently determining whether to decompress cached files or obtain uncompressed versions of those cached files.

Referring also to FIG. 7, data caching process 10 may be configured to receive 350 read request 120 on first cache system 126, wherein read request 120 identifies previously-written content (e.g., content 132) included within data array 130. Data caching process 10 may determine 352 if content 132 (identified in read request 120) is locally stored within first cache system 126 in a compressed format (e.g., as compressed content 132), and if content 132 (identified in read request 120) is stored within data array 130 in an uncompressed format (e.g., as content 132).

If it is determined 352 that the previously-written content requested in read request 120 is available locally (from first cache system 126) in a compressed format (e.g., compressed content 134) and is available remotely (from data array 130) in an uncompressed format (e.g., content 132), data caching process 10 may determine 354 if it is less computationally expensive to obtain previously-written content requested in read request 120 from data array 130 in an uncompressed format (e.g., content 132). Specifically, it may consume considerable processing power to decompress compressed content. Accordingly, when a compressed copy (e.g., compressed content 134) of the previously-written content requested in read request 120 is available locally (e.g. from first cache system 126) and an uncompressed copy (e.g., content 132) of the previously-written content requested in read request 120 is available remotely (e.g. from data array 130), it may consume less processing power to obtain the uncompressed copy of the previously-written content requested in read request 120 from data array 130 then it would consume to decompress the compressed copy of the previously-written content requested in read request 120 from first cache system 126.

There are many factors that data caching process 10 may take into consideration when making determination 354. For example, the compression ratio (as discussed above) of the compressed content may be taken into consideration when making determination 354, as content having higher compression ratios may be more computationally expensive to decompress than content having lower compression ratios. Additionally, the current level of network traffic between data array 130 and server computer/controller 100 may be taken into consideration, as additional loading of a crowded network may be deemed undesirable by data caching process 10. Further, the level of loading of the processor(s) included within server computer/controller 100 may be taken into consideration by data caching process 10, as additional loading of the processor(s) included within server computer/controller 100 may be deemed undesirable by data caching process 10. Additionally, the level of loading of the processor(s) included within data array 130 may be taken into consideration by data caching process 10, as additional loading of the processor(s) included within data array 130 may be deemed undesirable by data caching process 10.

12

If data caching process 10 determines 354 that it is less computationally expensive to obtain the previously-written content requested in read request 120 in an uncompressed format (e.g., content 132) from data array 130, data caching process 10 may obtain 356 the uncompressed content (e.g., content 132) stored within data array 130.

If data caching process 10 determines 354 that it is not less computationally expensive to obtain the previously-written content requested in read request 120 in an uncompressed format (e.g., content 132) from data array 130, data caching process 10 may decompress 358 the compressed content (e.g., compressed content 134) stored within first cache system 126.

As discussed above, data array 130 may include second cache system 128. Accordingly, when data caching process 10 determines 352 that the previously-written content requested in read request 120 is available locally (from first cache system 126) in a compressed format (e.g., compressed content 134) and is available remotely (from data array 130) in an uncompressed format (e.g., content 132), data caching process 10 may also/alternatively determine 360 if the previously-written content identified in read request 120 is stored within second cache system 128 of data array 130 in an uncompressed format.

Backend Content Copy:

Data caching process 10 may be configured to minimize/reduce network loading by intelligently determining whether data located in data array 130 should be copied from a first location to a second location within data array 130 (as opposed to being transmitted over network infrastructure 114).

Referring also to FIG. 8, data caching process 10 may be configured to receive 400 write request 116 on first cache system 126, wherein write request 116 identifies new content (e.g., content 118) to be written to data array 130.

As discussed above, data write request 116 may include a storage address 200 that defines the intended storage location within storage array 130 at which content 118 is to be stored. For example, storage address 200 may define a particular logical unit within data array 130 (e.g., a LUN or Logical Unit Number) and a particular storage address within that specific logical unit (e.g., an LBA or Logical Block Address) for storing content 118.

As discussed above, a content identifier (e.g., content identifier 256) may be a mathematical representation of a specific piece of content that may allow e.g. server computer/controller 100 to quickly determine whether two pieces of content are identical, as identical pieces of content would have identical content identifiers.

In one particular embodiment, content identifier 256 may be a hash function of the content. Upon receiving 400 write request 116 and content 118, data caching process 10 may generate 402 content identifier 256 for content 118. As discussed above, content identifier 256 generated 402 for the content (i.e., content 118) identified within write request 116 may be a hash function (e.g., a cryptographic hash) of content 118.

As is known in the art, a hash function is an algorithm/subroutine that maps large data sets to smaller data sets. The values returned by a hash function are typically called hash values, hash codes, hash sums, checksums or simply hashes. Hash functions are mostly used to accelerate table lookup or data comparison tasks such as e.g., finding items in a database and detecting duplicated or similar records in a large file.

For example, assume for illustrative purposes that write request 116 includes storage address 200 that defines the intended storage location for content 118 as LUN0/LBA0.

13

Accordingly, upon receiving **400** write request **116**, data caching process **10** may generate **402** content identifier **256** for content **118**. Assume for illustrative purposes that data caching process **10** generates a hash of content **118**, resulting in the generation **402** of content identifier **256** (namely hash value **acdfcla**).

This newly-generated content identifier **256** (i.e. **acdfcla**) associated with content **118** may be compared **404** to each of the other content identifiers (namely **abalabz**, **alazcha**, **abalabz**, **alazcha**) included within content directory **250** for first cache system **126** to determine if the newly-generated content identifier **256** (i.e. **acdfcla**) matches any of the other content identifiers (namely **abalabz**, **alazcha**, **abalabz**, **alazcha**) included within content directory **250**.

As discussed above, each entry of the plurality of entries **252** included within content directory **250** is associated with a unique piece of content included within first cache system **126**. Accordingly, each unique content identifier included within content directory **250** may be associated with a unique piece of content written to first cache system **126**.

If, when performing comparison **404**, data caching process **10** does not identify a content identifier (i.e., **abalabz**, **alazcha**, **abalabz**, **alazcha**) within content directory **250** that matches the above-described, newly-generated content identifier (i.e. **acdfcla**), data caching process **10** may write **406** content **118** to first cache system **126** and may provide **408** a copy of content **118** to data array **130** for storage within data array **130**. Additionally, data caching process **10** may modify **410** content directory **250** to include a new entry (i.e., entry **258**) that defines the newly-generated content identifier (i.e. **acdfcla**), the location of content **118** within first cache system **126** (i.e., **001011**), and the location of content **118** within data array **130** (i.e., **LUN0/LBA0**).

If, when performing comparison **404**, data caching process **10** identified a content identifier within content directory **250** that matched the above-described, newly-generated content identifier (i.e. **acdfcla**), data caching process **10** would perform differently.

Accordingly, further assume for illustrative purposes that a second write request (i.e., write request **116'**) includes storage address **200'** that defines the intended storage location for content **118'** as **LUN0/LBA1**. Accordingly, upon receiving **400** write request **116'**, data caching process **10** may generate **402** content identifier **256** for content **118'**. Assume for illustrative purposes that data caching process **10** generates a hash of content **118'**, resulting in the generation **402** of content identifier **256** (namely hash value **alazcha**).

This newly-generated content identifier **256** (i.e. **alazcha**) associated with content **118'** may be compared **404** to each of the other content identifiers (namely **abalabz**, **alazcha**, **abalabz**, **alazcha**) included within content directory **250** for first cache system **126** to determine if the newly-generated content identifier **256** (i.e. **alazcha**) matches any of the other content identifiers (namely **abalabz**, **alazcha**, **abalabz**, **alazcha**) included within content directory **250**.

If, when performing comparison **404**, data caching process **10** does identify a content identifier (namely **alazcha**) within content directory **250** that matches the above-described, newly-generated content identifier (i.e. **alazcha**), data caching process **10** may perform a couple of functions.

For example, data caching process **10** may copy **412** content on data array **130** from a first location (namely **LUN4/LBA7**) on data array **130** associated with the matching content identifier to a second location (namely **LUN0/LBA1**) on data array **130** associated with the newly-generated content identifier. Specifically and for the reasons discussed above, since data caching process **10** identified matching content

14

identifiers (i.e., having the same value of **alazcha**), content **118'** must be identical to the content stored at **LUN4/LBA7**. Accordingly, instead of transmitting content **118'** from server computer/controller **100** to data array **130**, data caching process **10** may simply effectuate on data array **130** a copy of the data currently stored at **LUN4/LBA7** to **LUN0/LBA1**, thus freeing up server computer/controller **100** to perform other functions and reducing network traffic on network infrastructure **114**.

Further, data caching process **10** may modify **414** the entry (i.e., entry **260**) within content directory **250** that is associated with the matching content identifier (i.e., **alazcha**) to include storage address **200'** that defines the intended storage location for content **118'** (i.e., **LUN0/LBA1** within data array **130**), thus generating modified entry **260'**. Accordingly, modified entry **260** identifies that the pieces of content that are currently stored at **LUN4/LBA7** and **LUN0/LBA1** within data array **130** are identical. Accordingly, a single piece of cached content (located first cache address **010111** within first cache system **126**) may be used as a local cached copy for both pieces of content stored on data array **130**.

Content Identifier Backend Retrieval:

Data caching process **10** may be configured to minimize/reduce loading of server computer/controller **100** by determining if content already exists on first cache system **126** before retrieving the same from data array **130**.

Referring also to FIG. 9, data caching system **10** may be configured to receive **450** read request **120** on first cache system **126**, wherein read request **120** identifies previously-written content (e.g., content **132**) included within data array **130**. Upon receiving **450** read request **120**, data caching process **10** may process read request **120** to determine if content **132** is locally available within first cache system **126**.

As discussed above, read request **120** may define a particular logical unit within data array **130** (e.g., a **LUN** or **Logical Unit Number**) and a particular storage address within that specific logical unit (e.g., an **LBA** or **Logical Block Address**) for retrieving content **132** sought from data array **130**.

As discussed above, content directory **250** may include plurality of entries **252**, wherein each of these entries may identify: data array storage address **200/202** (e.g. a logical storage unit and a storage address at which a specific piece of previously-written content is located within storage array **130**); first cache address **254** (e.g., the location within first cache system **126** at which the specific piece of previously-written content is also located), and content identifier **256** for the specific piece of previously-written content.

Accordingly, when data caching process **10** processes read request **120** to determine if content **132** is available within first cache system **126**, data caching process **10** may review content directory **250** to see if the **LUN** and **LBA** defined within read request **120** is listed within content directory **250**.

Assume for illustrative purposes that read request **120** is requesting the data stored at **LUN6/LBA4** of data array **130**. As **LUN6/LBA4** is not listed as a data array storage address **200/202** within content directory **250**, data caching process **10** will determine that the content requested in read request **120** is not available locally (i.e. from first cache system **126**).

Accordingly, data caching process **10** may request **452**, from the data array, the content identifier for the content identified in the read request (namely content **132**). Upon receiving request **452** from data caching process **10**, data array **130** may either a) retrieve the content identifier for the content identified in read request **120** (if previously generated and stored within e.g. a content directory (not shown) for second cache system **128**; or b) generate the content identifier for the content identified in the read request **120** (as data array

15

130 has access to content 132). Once generated/retrieved by data array 130, content identifier 138 may be received 454 by data caching process 10 from data array 130. As discussed above, content identifier 138 may be a hash function of the content requested in data request 120 (namely content 132).

Data caching process 10 may compare 456 content identifier 138 (i.e., the content identifier associated with the content identified within read request 120, namely content 132) to each of the content identifiers (namely abalabz, alazcha, abalabz, alazcha) included within content directory 250 for first cache system 126 to determine if a matching content identifier exists. As discussed above, each of the plurality of content identifiers (namely abalabz, alazcha, abalabz, alazcha) included within content directory 250 is associated with a piece of content previously-written to first cache system 126.

If, when performing comparison 456, data caching process 10 does not identify a content identifier included within content directory 250 that matches content identifier 138, data caching process 10 may obtain 458 the previously-written content identified in read request 120 (i.e., content 132) from data array 130 and may write 460 content 132 to first cache system 126 (at e.g., first cache address 111011). Additionally, data caching process 10 may modify 462 content directory 250 to include a new entry (i.e., entry 262) that defines content identifier 138 (namely ablcboa), the location of content 132 within first cache system 126 (i.e., first cache address 111011), and the location of content 132 within data array 130 (i.e., LUN6/LBA4).

If, when performing comparison 456, data caching process 10 identifies a content identifier (either abalabz, alazcha, abalabz or alazcha) included within content directory 250 that matches content identifier 138, data caching process 10 may retrieve 464 content 132 identified in read request 120 from first cache system 126.

Further, data caching process 10 may also modify 466 an entry within content directory 250 that is associated with the matching content identifier to include a storage address that defines the storage location of content 132 (i.e., LUN6/LBA4 within data array 130). This modification of a data entry within content directory 250 would be similar to the manner in which entry 260 was modified to create modified entry 260'.

For example, assume that the matching content identifier was again alazcha which was included within entry 260. Accordingly, entry 260 would be modified to also identify the content stored at LUN6/LBA4 (i.e., content 132) in addition to the content stored at LUN4/LBA7, as the content at LUN6/LBA4 and LUN4/LBA7 within data array 130 are identical. Accordingly, a single piece of cached content within first cache system 126 may be used as a local cached copy for both pieces of content stored on data array 130.

#### Content Identifier Generation

Data caching process 10 may be configured to conserve storage space on data array 130 by deleting redundant data.

Referring also to FIG. 10, data caching process 10 may be configured to receive 500 read request 120 on first cache system 126, wherein read request 120 identifies previously-written content (e.g., content 132) included within data array 130. Upon receiving 450 read request 120, data caching process 10 may process read request 120 to determine if content 132 is locally available within first cache system 126.

As discussed above, read request 120 may define a particular logical unit within data array 130 (e.g., a LUN or Logical Unit Number) and a particular storage address within that specific logical unit (e.g., an LBA or Logical Block Address) for retrieving content 132 sought from data array 130.

16

As discussed above, content directory 250 may include plurality of entries 252, wherein each of these entries may identify: data array storage address 200/202 (e.g. a logical storage unit and a storage address at which a specific piece of previously-written content is located within storage array 130); first cache address 254 (e.g., the location within first cache system 126 at which the specific piece of previously-written content is also located), and content identifier 256 for the specific piece of previously-written content.

Accordingly, when data caching process 10 processes read request 120 to determine if content 132 is available within first cache system 126, data caching process may review content directory 250 to see if the LUN and LBA defined within read request 120 is listed within content directory 250.

Assume again for illustrative purposes that read request 120 is requesting the data stored at LUN6/LBA4 of data array 130. As LUN6/LBA4 is not listed as a data array storage address 200/202 within content directory 250, data caching process 10 will determine that the content requested in read request 120 is not available locally (i.e. from first cache system 126).

Accordingly, data caching process 10 may need to obtain 502 content 132 identified in read request 120 from data array 130. In this particular example, since LUN6/LBA4 is not defined within content directory 250, a local cached copy of content 132 is not present/available within first cache system 126 and data caching process 10 will need to obtain 502 content 132 from data array 130.

Once content 132 is obtained 502 from data array 130, data caching process 10 may generate 504 a content identifier (e.g., content identifier 138) for content 132. As discussed above, content identifier 138 generated 504 for the content 132 identified within read request 120 may be a hash function of content 132. Assume for illustrative purposes that data caching process 10 generates a hash of content 132, resulting in the generation 504 of content identifier 138 (namely hash value ablcboa).

Data caching process 10 may compare 506 content identifier 138 (i.e., the content identifier associated with the content identified within read request 120, namely content 132) to each of the content identifiers (e.g., abalabz, alazcha, abalabz, alazcha) included within content directory 250 for first cache system 126 to determine if a matching content identifier exists. As discussed above, each of the plurality of content identifiers (namely abalabz, alazcha, abalabz, alazcha) included within content directory 250 is associated with a piece of content previously-written to first cache system 126.

If, when performing comparison 506, data caching process 10 does not identify a content identifier included within content directory 250 that matches content identifier 138, data caching process 10 may write 508 content 132 to first cache system 126 (at e.g., first cache address 111011). Additionally, data caching process 10 may modify 510 content directory 250 to include a new entry (i.e., entry 262) that defines content identifier 138 (namely ablcboa), the location of content 132 within first cache system 126 (i.e., first cache address 111011), and the location of content 132 within data array 130 (i.e., LUN6/LBA4).

If, when performing comparison 456, data caching process 10 identifies a content identifier (either abalabz, alazcha, abalabz or alazcha) included within content directory 250 that matches content identifier 138, data caching process 10 may modify 512 an entry within content directory 250 that is associated with the matching content identifier to include a storage address that defines the storage location of content 132 (i.e., LUN6/LBA4 within data array 130). This modifi-

17

cation of a data entry within content directory **250** would be similar to the manner in which entry **260** was modified to create modified entry **260'**.

Assume for illustrative purposes that data caching process **10** generates a hash of content **132**, resulting in the generation **504** of content identifier **138** having a hash value of ablcboa. As the matching content identifier is included within entry **260** of content directory **250**, entry **260** may be modified **512** by data caching process **10** to also identify the piece of content stored at LUN6/LBA4 (i.e., content **132**) in addition to the piece of content stored at LUN4/LBA7, as the content at LUN6/LBA4 and LUN4/LBA7 within data array **130** are identical.

Additionally, if a matching content identifier is identified, data caching process **10** may notify **514** data array **130** that content **130** identified in read request **120** is identical to the previously-written content associated with the matching content identifier. Accordingly and in this example, content **130** stored at LUN6/LBA4 within data array **130** is identical to the content stored at LUN4/LBA7 within data array **130**. Once notified **516**, data array **130** may e.g., delete one of the redundant copies (e.g., at LUN6/LBA4 or LUN4/LBA7) on e.g., non-volatile, electro-mechanical memory system **112** and/or second cache system **128** and may map LUN6/LBA4 and LUN4/LBA7 to a single memory location within e.g., non-volatile, electro-mechanical memory system **112** and/or second cache system **128**.

As will be appreciated by one skilled in the art, the present disclosure may be embodied as a method, system, or computer program product. Accordingly, the present disclosure may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, the present disclosure may take the form of a computer program product on a computer-usable storage medium having computer-usable program code embodied in the medium.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present disclosure. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the disclosure. As used herein, the singular forms "a", "an" and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms "comprises" and/or "comprising," when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/

18

or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of the present disclosure has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the disclosure in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the disclosure. The embodiment was chosen and described in order to best explain the principles of the disclosure and the practical application, and to enable others of ordinary skill in the art to understand the disclosure for various embodiments with various modifications as are suited to the particular use contemplated.

A number of implementations have been described. Having thus described the disclosure of the present application in detail and by reference to embodiments thereof, it will be apparent that modifications and variations are possible without departing from the scope of the disclosure defined in the appended claims.

What is claimed is:

1. A computer-implemented method of processing read requests comprising:
  - receiving a read request on a first cache system, wherein the read request identifies previously-written content included within a data array;
  - obtaining the previously-written content identified in the read request from the data array;
  - generating a read request content identifier for the previously-written content identified in the read request;
  - comparing the read request content identifier associated with the previously-written content identified in the read request to a plurality of content identifiers included within a content directory for the first cache system to determine if a matching content identifier exists, wherein each of the plurality of content identifiers is associated with a piece of previously-written content included within the first cache system; and
  - deleting at least one redundant copy of the previously-written content.
2. The computer-implemented method of claim 1 further comprising:
  - if a matching content identifier is not identified, writing the previously-written content identified in the read request to the first cache system.
3. The computer-implemented method of claim 1 further comprising:
  - if a matching content identifier is identified, notifying the data array that the previously-written content identified in the read request is identical to previously-written content associated with the matching content identifier.
4. The computer-implemented method of claim 1 wherein the read request identifies a logical unit and a storage address of the previously-written content identified in the read request on the data array.
5. The computer-implemented method of claim 4 further comprising:
  - if a matching content identifier is identified, modifying an entry within the content directory that is associated with the matching content identifier to identify the logical

19

storage unit and the storage address of the previously-written content identified in the read request on the data array.

6. The computer-implemented method of claim 1 wherein the read request content identifier is a hash function of the previously-written content identified in the read request.

7. The computer-implemented method of claim 1 wherein the content directory for the first cache system includes a plurality of entries, wherein each of the plurality of entries is associated with a specific piece of previously-written content within the first cache system, wherein each of the plurality of entries within the content directory identifies:

- a logical storage unit and a storage address at which the specific piece of previously-written content is located within the storage array;
- a first cache address at which the specific piece of previously-written content is located within the first cache system, and
- a content identifier for the specific piece of previously-written content.

8. A computer program product residing on a non-transitory computer readable medium having a plurality of instructions stored thereon which, when executed by a processor, cause the processor to perform operations comprising:

- receiving a read request on a first cache system, wherein the read request identifies previously-written content included within a data array;
- obtaining the previously-written content identified in the read request from the data array;
- generating a read request content identifier for the previously-written content identified in the read request;
- comparing the read request content identifier associated with the previously-written content identified in the read request to a plurality of content identifiers included within a content directory for the first cache system to determine if a matching content identifier exists, wherein each of the plurality of content identifiers is associated with a piece of previously-written content included within the first cache system; and
- if a matching content identifier is identified, deleting at least one redundant copy of the previously-written content.

9. The computer program product of claim 8 further comprising instructions for:

- if a matching content identifier is not identified, writing the previously-written content identified in the read request to the first cache system.

10. The computer program product of claim 8 further comprising instructions for:

- if a matching content identifier is identified, notifying the data array that the previously-written content identified in the read request is identical to previously-written content associated with the matching content identifier.

11. The computer program product of claim 8 wherein the read request identifies a logical unit and a storage address of the previously-written content identified in the read request on the data array.

12. The computer program product of claim 11 further comprising instructions for:

- if a matching content identifier is identified, modifying an entry within the content directory that is associated with the matching content identifier to identify the logical storage unit and the storage address of the previously-written content identified in the read request on the data array.

20

13. The computer program product of claim 8 wherein the read request content identifier is a hash function of the previously-written content identified in the read request.

14. The computer program product of claim 8 wherein the content directory for the first cache system includes a plurality of entries, wherein each of the plurality of entries is associated with a specific piece of previously-written content within the first cache system, wherein each of the plurality of entries within the content directory identifies:

- a logical storage unit and a storage address at which the specific piece of previously-written content is located within the storage array;
- a first cache address at which the specific piece of previously-written content is located within the first cache system, and
- a content identifier for the specific piece of previously-written content.

15. A computing system including at least one processor and at least one memory architecture coupled with the at least one processor, wherein the computing system is configured to perform operations comprising:

- receiving a read request on a first cache system, wherein the read request identifies previously-written content included within a data array;
- obtaining the previously-written content identified in the read request from the data array;
- generating a read request content identifier for the previously-written content identified in the read request;
- comparing the read request content identifier associated with the previously-written content identified in the read request to a plurality of content identifiers included within a content directory for the first cache system to determine if a matching content identifier exists, wherein each of the plurality of content identifiers is associated with a piece of previously-written content included within the first cache system; and
- deleting at least one redundant copy of the previously-written content.

16. The computing system of claim 15 further configured to perform operations comprising:

- if a matching content identifier is not identified, writing the previously-written content identified in the read request to the first cache system.

17. The computing system of claim 15 further configured to perform operations comprising:

- if a matching content identifier is identified, notifying the data array that the previously-written content identified in the read request is identical to previously-written content associated with the matching content identifier.

18. The computing system of claim 15 wherein the read request identifies a logical unit and a storage address of the previously-written content identified in the read request on the data array.

19. The computing system of claim 18 further configured to perform operations comprising:

- if a matching content identifier is identified, modifying an entry within the content directory that is associated with the matching content identifier to identify the logical storage unit and the storage address of the previously-written content identified in the read request on the data array.

20. The computing system of claim 15 wherein the read request content identifier is a hash function of the previously-written content identified in the read request.

21. The computing system of claim 15 wherein the content directory for the first cache system includes a plurality of entries, wherein each of the plurality of entries is associated

**21**

with a specific piece of previously-written content within the first cache system, wherein each of the plurality of entries within the content directory identifies:

a logical storage unit and a storage address at which the specific piece of previously-written content is located 5 within the storage array;

a first cache address at which the specific piece of previously-written content is located within the first cache system, and

a content identifier for the specific piece of previously-written 10 content.

\* \* \* \* \*

**22**